

ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, Illinois 60439

ANL/CGC-005-0401

**Computational Design and Performance of the
Fast Ocean Atmosphere Model, Version One***

by

Robert Jacob,¹ Chad Schafer,² Ian Foster,¹ Michael Tobis,³ and John Anderson³

¹ Argonne National Laboratory, Mathematics and Computer Science Division,
9700 S. Cass Ave., Argonne, IL 60439, USA {jacob,itf}@mcs.anl.gov

² University of California at Berkeley, Department of Statistics,
367 Evans Hall, Berkeley, CA 94720, USA

³ University of Wisconsin-Madison, Space Science and Engineering Center,
1225 W. Dayton St., Madison, WI 53706, USA

Climate and Global Change

Series No. ANL/CGC-005-0401

April 2001

*This work was supported by the Office of Biological and Environmental Research of the U.S. Department of Energy's Office of Science under Contracts W-31-109-ENG-38 and DE-FG02-98ER62617.
© Springer-Verlag

Contents

Abstract	1
1 Introduction	1
2 Component Models	2
3 Design of the Coupled System	3
3.1 Number of Numerical Grids	3
3.2 Decomposition of Component Models	3
3.3 Flux Interpolation and Calculation	4
3.4 FOAM's Coupler	6
3.5 Constructing the Full Model	6
4 Model Performance	6
5 Conclusions	8
References	9

Computational Design and Performance of the Fast Ocean Atmosphere Model, Version One

by

Robert Jacob, Chad Schafer, Ian Foster, Michael Tobis, and John Anderson

Abstract

The Fast Ocean Atmosphere Model (FOAM) is a climate system model intended for application to climate science questions that require long simulations. FOAM is a distributed-memory parallel climate model consisting of parallel general circulation models of the atmosphere and ocean with complete physics parameterizations as well as sea-ice, land surface, and river transport models. FOAM's coupling strategy was chosen for high throughput (simulated years per day). A new coupler was written for FOAM and some modifications were required of the component models. Performance data for FOAM on the IBM SP3 and SGI Origin2000 demonstrates that it can simulate over thirty years per day on modest numbers of processors.

1 Introduction

The Earth's climate is the long-term average of the behavior of the ocean, land surface, and sea ice. These components exchange heat, momentum, and fresh water with each other through their common interfaces. The difference in heat storage times, fluid dynamic scales, and transport/storage of water results in complex interactions between the components.

To study the Earth's climate, researchers have developed coupled models, comprising an atmospheric general circulation model (GCM), a model of the ocean general circulation, a model of the dynamic and thermodynamic properties of sea ice, and a model of the temperature and composition of the land surface. Typically, each component model is a separately developed program with its own code style, language, data structures, numerical methods, and associated numerical grids. Constructing a single model from these submodels requires several design decisions about grid structure, model decomposition, and flux calculations. Additionally, a "coupler" must be created to enable the subcomponents to exchange heat and other quantities.

While many coupled climate models are used in studies that require predicting fine-scale details of temperature and precipitation changes, other climate science studies are more concerned with high throughput than with high resolution. For example, paleoclimate scientists wish to know how the climate responded to conditions thousands of years ago, when the earth received more solar radiation in the summer or when glaciers covered much of the northern hemisphere; and climate scientists wish to know how much low-frequency internal variability the climate system has and how it might interact with anthropogenic change. Integrating a climate model to even a quasi-equilibrium for the myriad of interesting paleoclimate scenarios requires hundreds of simulated years, while studying low-frequency variability requires simulations of thousands of years. Thus there is a compelling need for a climate model with high-throughput—a high number of simulated years per day.

The Fast Ocean Atmosphere Model is a coupled climate model designed to meet this need. It has already been used to examine low-frequency variability and paleoclimate problems [7, 11, 13]. This paper describes how the goal of high-throughput guided the design of

FOAM and its coupler. Section 2 describes the component models of FOAM. The unique elements of FOAM’s software design are described in Section 3. In Section 4, performance data for FOAM is presented in the form of throughput per processor for various platforms. Section 5 briefly outlines future work on the next version of FOAM.

2 Component Models

The first design decision made for FOAM was to use distributed-memory parallel components. The low cost/performance ratios of distributed-memory multicomputers make them the platform of choice for the kind of climate science questions FOAM is intended to address. Fortunately, FOAM was able to use several existing parallel component models, as briefly described below. (A detailed description of the physical equations solved by each component is beyond the scope of this paper; however, that information can be found in the cited references.)

The atmospheric component of FOAM is derived from PCCM2, developed by Argonne National Laboratory, Oak Ridge National Laboratory, and the National Center for Atmospheric Research. PCCM2 is a functionally equivalent version of NCAR’s CCM2 in which the physics and dynamics calculations have been given a two-dimensional parallel decomposition (see Figure 1a for an example). The basic equations and numerical methods used

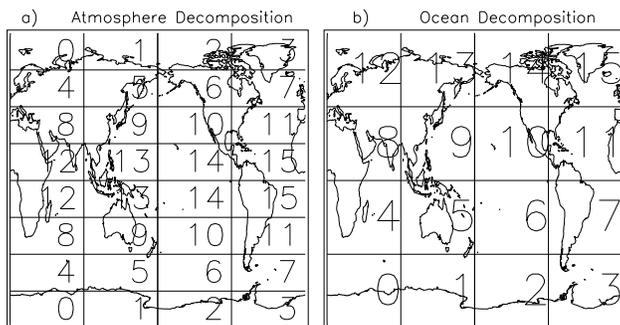


Figure 1: A depiction of how physical space is partitioned among 4x4 processors according to the decomposition strategy of (a) the atmosphere and (b) the ocean.

in CCM2 (Eulerian-spectral elements for the dry dynamics and semi-Lagrangian transport for moisture) are described by Hack et al. [6], and the alterations necessary to implement a parallel version and its performance on various parallel platforms are described by Drake et al. [5]. When CCM3 [8] was released, the FOAM development team added the new physics parameterizations to the atmosphere model to create the current atmosphere component of FOAM, informally called PCCM3-UW.

The ocean component of FOAM, called OM3, is also a parallel model with a two-dimensional decomposition (Fig. 1b). OM3 is a z-coordinate finite-difference model. While documentation specific to OM3 is still under development, the basic equations solved are the same as those for the widely used z-coordinate Modular Ocean Model created by GFDL [4] [3]. However, FOAM uses a free surface as described by Killworth et al. [9] and contains

numerical methods specifically chosen for their efficiency on distributed-memory parallel platforms [14].

The land surface and sea-ice models in version one of FOAM are based on those of PCCM2 [6] but with some important modifications. The prescribed evaporation and snow cover have been replaced with a prognostic box hydrology model from CCM1 [15]. Also, the effect of sea-ice creation/destruction on ocean salinity has been included.

Two new software components had to be created to complete FOAM. The coupler will be described below. The other new component is a parallel river transport model. In order to prevent a constant increase in salinity of the ocean in long climate simulations, it was necessary to close the hydrologic cycle by returning precipitation that falls on land to the ocean. FOAM's river model is a parallel implementation of concepts described by Miller et al. [12].

FOAM and its individual components use the Message Passing Interface (MPI) library for communication. The land and sea-ice models do not require any communication, however, since they are implemented as one-dimensional models at each land and ocean point on the surface, respectively.

3 Design of the Coupled System

In parallel climate models, each component, especially the atmosphere, can place large demands on the communication network and CPU resources of modern parallel supercomputers. A coupled system creates additional demands on the bandwidth (mostly from the exchange of data between models) and on CPU resources (mostly from the interpolation of data between numerical grids). Design choices in FOAM were made to minimize the impact of coupling. Those choices and the reasons for them are summarized in this section.

3.1 Number of Numerical Grids

Generally, each model may have its own finite difference grid. FOAM's design starts with the requirement that all surface models (land, sea ice, and ocean) use the grid of the ocean model. The intention is to eliminate a class of operations in the coupler that would be required to keep track of land/ocean fraction and merge the calculated fluxes accordingly. This decision does not have a large impact on the performance of the model, but it does greatly simplify the conceptual picture of the coupled system. It also has implications for other design choices described below.

3.2 Decomposition of Component Models

Just as each model may have its own numerical grid, each parallel model may decompose that grid in a different way over its set of MPI processes. Additionally, each model may have physically separate pools of processors. In order to calculate a new flux or provide to a model an internally calculated flux, such as precipitation, data that occupies the same physical space must be brought into the same physical piece of memory. The bandwidth cost of this communication will depend on how dissimilar the decompositions are and how much overlap there is of model state in memory. To minimize this cost, FOAM restricts the

possible decomposition of some of the component models based on a consideration of the time scales between components, as shown in Figure 2. In general, an atmosphere model

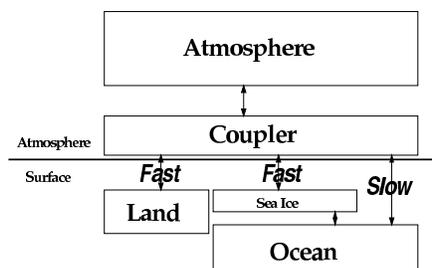


Figure 2: Schematic of the components of FOAM indicating which components interact on fast and slow time steps and reflecting the physical relationship between components. The parallel river model is omitted for clarity.

has a shorter basic time step than does the ocean because of differences in the dominant fluid dynamical scales in the flows. In FOAM, the atmosphere has a basic time step of 30 minutes, while the ocean uses a 6-hour time step. Because of the small heat capacity of land and sea ice, their temperature structure varies with the diurnal cycle of solar heating (which is resolved by the atmosphere model), and thus they need to communicate with the atmosphere every one or two time steps. The ocean, however, with a much larger heat capacity, does not have a significant diurnal cycle and needs to communicate with the atmosphere only once a day.

In order to reduce the costs of communicating fluxes, the “fast” components (Fig. 2) are all given the same decomposition so that there is nearly a one-to-one correspondence in area coverage between atmosphere, land, and sea ice. The atmosphere, land, and sea-ice models are executed sequentially on the same pool of processors so nearly all data necessary for physical coupling resides in local memory. The area coverage will not match exactly, however, because the land and sea ice are on the ocean’s numerical grid, which coincides nowhere with the atmosphere’s. Some communication is required to update points on the edges of the local rectangular regions covered by each MPI process, but this communication overhead is less than what would be required if the fast components had to transfer or transpose their entire state across multiple processors every one or two time steps.

The ocean model has its own decomposition and is assigned a distinct set of processors. It integrates concurrently with the atmosphere, land, and sea ice. Trial and error determines how many processors to allocate to each side to avoid blocking on data. Communication between the ocean and atmosphere–land–sea-ice coupler is done by designating one MPI process on each side to receive all the data from the other component and redistribute it according to that component’s decomposition strategy. Since the ocean is the slow component and needs to communicate less frequently, the impact of this single-node bottleneck is minimized.

3.3 Flux Interpolation and Calculation

One of the basic tasks of a coupled model is to interpolate quantities from one grid to another. Calculating and interpolating fluxes in FOAM are simplified by placing the land,

sea-ice, and ocean models all onto a common grid as mentioned above. Thus, the number of grids in the model is reduced to at most two instead of possibly one per model.

FOAM considers the surface of the globe as being divided into rectangles or tiles. The center of each tile is at a computational point of the model, and the four edges are halfway between each of the neighboring computational points. A third set of tiles is made by laying the ocean grid on top of the atmosphere grid, as in Fig. 3a. New fluxes are calculated on this overlap grid set of tiles. The flux through a model tile is constructed by area-weighted averaging the appropriate pieces of the overlap tiles as shown in Fig. 3b. Global conservation is assured as long as all the overlap pieces are used once and only once in the averaging calculation for each grid.

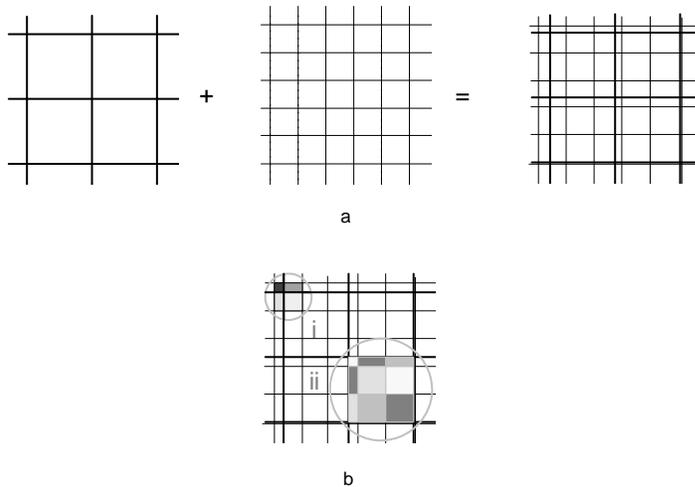


Figure 3: (a) Forming the overlap grid. (b) A surface (region i) and atmosphere tile (region ii) composed of overlap pieces.

The interpolation method can be demonstrated further by considering the calculation of sensible heat. This requires the temperature of the lowest atmosphere level TA and a surface temperature TS . FOAM's calculation of the sensible heat flux through an overlap tile can be approximated as

$$SH = K(TS_{i,j,k} - TA_{l,m}), \quad (1)$$

where K is a constant. The triplet i, j, k indexes the pieces a given surface tile i, j is divided into (1, 2, or 4 pieces), and l, m indexes the atmosphere tile lying above the piece. The mapping of the overlap tile indices onto each of the two grids and the area of the overlap tiles are calculated offline and stored in lookup tables for use during runtime. Once the flux at each overlap tile is known, an area-weighted aggregate is formed for each ocean and atmosphere tile in the subdomain of the coupler, as illustrated in Fig. 3b.

3.4 FOAM's Coupler

After all the simplifications made with component models, FOAM's coupler is relatively straightforward. The coupler is implemented as a subroutine of the atmosphere. The coupler performs all the communication necessary to resolve the incomplete overlapping of the local surface and atmosphere grids and calls the land, sea-ice, and river models (which are themselves subroutines of the coupler). The merging of fluxes coming from separate ocean, land, and sea-ice models into a single atmospheric cell is greatly simplified by the use of single surface grid. Finally, the coupler accumulates ocean fluxes and communicates with the ocean as necessary. "Flux corrections", which are sometimes needed to achieve a stable coupled climate, are not used in FOAM.

3.5 Constructing the Full Model

The noncoincidence of the two grids presented a difficulty when combining the coupler and fast surface models with the atmosphere. In the original atmosphere model, PCCM2, all physics routines were called once for each latitude, with the surface physics occurring about two thirds of the way through the sequence. The uneven overlap of the two grids means the coupler occasionally needs two atmosphere latitudes to calculate the flux through one ocean latitude. Consequently, the calling tree for the PCCM2 physics was split so that all the physics before the surface package could be completed for the whole globe before calling the coupler and surface models. The atmosphere resumes execution when the coupler returns. This split is still present in FOAM's current atmosphere model PCCM3-UW. (NCAR's CCM3, which unlike PCCM2 was designed to be part of a coupled system, contains a similar split [1].)

FOAM is a single executable image. A small main program divides processors between the atmosphere-coupler-surface component and the ocean according to compile-time settings and then calls each. Some minor modifications to the atmosphere and ocean were necessary to turn them into subroutines. A summary of FOAM's structure is shown in Fig. 4.

The atmosphere and ocean are each told by the main routine how many days to integrate. The atmosphere directs the execution of itself, the coupler, and the surface components that share the same processors. The ocean model waits to receive fluxes from the coupler and send sea surface temperatures according to a frequency set at compile time (currently every ocean time step).

4 Model Performance

The performance of FOAM and some of its components is presented in Figure 5. While the atmosphere model supports many choices for truncation level of the spectral transform numerical method, the standard configuration for FOAM uses a rhomboidal truncation with fifteen wave numbers. The associated physical grid has 40 latitudes, 48 longitudes, and 18 levels. The standard resolution for the ocean model is a Mercator grid with 128 latitudes, 128 longitudes, and 16 levels. Runs were conducted on three different platforms using configurations of 5 processors (2 by 2 for the atmosphere-land-sea-ice-coupler and 1 for the ocean), 9 processors (2 by 4 plus 1), 17 processors (4 by 4 plus 1), 34 (4 by 8

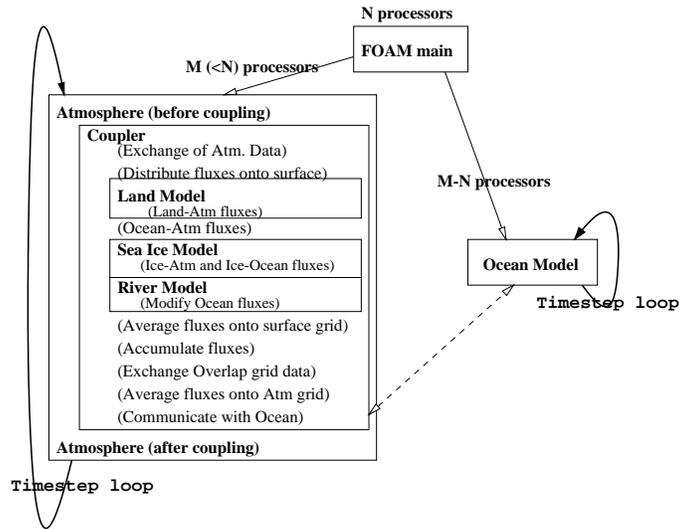


Figure 4: A schematic of FOAM's software design.

plus 1 by 2), and 68 processors (4 by 16 plus 2 by 2) processors. (The last configuration comes from atmosphere model requirements for number of local latitudes and number of processors.) Figure 5 shows that the model scales well over the range of configurations tested. The throughput goal for FOAM was chosen to be ten thousand times real time, that is, the ability to simulate 10,000 years in a single calendar year. FOAM meets this goal on a modest number of processors.

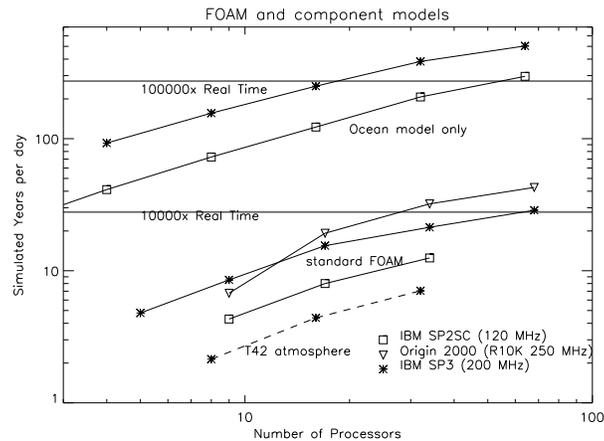


Figure 5: Timings for FOAM and its components on various platforms for different numbers of processors.

Figure 5 and the processor configurations show how the atmosphere component dominates performance. The ocean model, despite having nearly eight times as many grid points as the atmosphere, has a much higher throughput when running by itself on the same number of processors. The difference comes from the numerous, detailed physical

parameterizations in the atmosphere compared with the ocean and other components. This behavior is seen in other coupled climate models also.

Given the dominance of the atmosphere component on FOAM's performance, the choice of atmosphere resolution is as important for overall throughput as the design of the coupler and the complete model. When the atmosphere model is executed by itself at a higher, more standard resolution of T42 (64 latitudes and 128 longitudes), it is slower by a factor roughly equal to the ratio of the number of points in each grid.

Figure 5 also shows how increases in processor speed have translated into increases in model throughput. Performance nearly doubles between the IBM SP2 and IBM SP3 and nearly doubles again when timed on an IBM SP3 with 350MHz CPU's (not shown).

5 Conclusions

The Fast Ocean Atmosphere Model is a coupled climate model created to address climate science questions that require many simulated years of interaction but not best-possible resolution. Through a combination of a low-resolution parallel atmosphere model, a highly efficient parallel ocean, and a software design that minimizes the effect of coupling on bandwidth and CPU usage, FOAM can simulate several decades of global climate interaction a day on modest numbers of distributed-memory parallel processors. The simulated climate is physically realistic and comparable to models with higher resolution.

The rapid development of FOAM was made possible by the work done by others on the atmosphere and ocean models and also by the simplicity of land and sea-ice components. Work on the next version of FOAM will concentrate on upgrading the representation of the sea-ice and land surface to match that of the NCAR Climate System Model (CSM) [2]. FOAM's current sea ice model lacks many standard features such as ice dynamics and complex thermodynamics. In order to accommodate new models, it may be necessary to relax the conditions FOAM imposes on the land and sea-ice components. Although the resolution of FOAM's ocean model is certainly adequate for modern sea-ice and land models, separately developed components do not usually have the flexibility to change their resolution, grid, or decomposition. A rewrite of FOAM's coupler will be necessary to support arbitrary decompositions of land and sea-ice models; fortunately, that task has been made much easier by the development of new software libraries such as the Model Coupling Toolkit [10].

The relationship between FOAM and larger climate modeling efforts such as CSM is complementary. FOAM uses CSM's submodels wherever possible; adapting them to FOAM's framework and applying them to climate questions that require more throughput than resolution. As large, institutionally supported climate models increase their flexibility and parallelism, it should be possible to match FOAM's combination of throughput and simulation quality by invoking appropriate options within the larger model. Until then, FOAM development will continue.

References

- [1] Thomas L. Acker, Lawrence E. Buja, James M. Rosinski, and John E. Truesdale. User's Guide to NCAR CCM3. NCAR Tech. Note NCAR/TN-421+IA, Natl. Cent. for Atmos. Res., Boulder, Co., 1996.
- [2] B. A. Boville and P. R. Gent. The NCAR climate system model, version one. *J. Climate*, 11:1115–1130, 1998.
- [3] Kirk Bryan. A numerical method for the study of the circulation of the World Ocean. *J. Comp. Phys.*, 4:347–376, 1969.
- [4] M. D. Cox. A Primitive Equation three-dimensional model of the ocean. Technical Report GFDL Ocean Group Tech. Rep. 1, GFDL, Princeton, NJ, 1984.
- [5] J. Drake, I. Foster, J. Michalakes, Brian Toonen, and Pat Worley. Design and performance of a scalable parallel community climate model. *Parallel Computing*, 21(10):1571–1591, October 1995.
- [6] James J. Hack, Byron A. Boville, Bruce P. Briegleb, Jeffrey T. Kiehl, Philip J. Rasch, and David L. Williamson. Description of the NCAR Community Climate Model (CCM2). NCAR Tech. Note NCAR/TN-382+STR, Natl. Cent. for Atmos. Res., Boulder, Co., 1993.
- [7] Robert Jacob. *Low Frequency Variability in a Simulated Atmosphere Ocean System*. PhD thesis, University of Wisconsin-Madison, 1997.
- [8] Jeffrey T. Kiehl, James J. Hack, Gordon B. Bonan, Byron A. Boville, Bruce P. Briegleb, David L. Williamson, and Philip J. Rasch. Description of the NCAR Community Climate Model (CCM3). NCAR Tech. Note NCAR/TN-420+STR, Natl. Cent. for Atmos. Res., Boulder, Co., 1996.
- [9] P. D. Killworth, D. Stainforth, D. J. Webb, and S. M. Patterson. The development of a free-surface Bryan-Cox-Semtner ocean model. *J. Phys. Oceanogr.*, 21:1333–1348, 1991.
- [10] Jay Larson and Robert Jacob. A message-passing parallel Model Coupling Toolkit (MCT). *in prep.*, 2001.
- [11] Zhengyu Liu, John Kutzbach, and Lixin Wu. Modeling climate shift of El-Niño variability in the holocene. *Geophys. Res. Lett.*, 27:2265–2268, 2000.
- [12] James R. Miller, Gary L. Russell, and Guilherme Caliri. Continental-scale river flow in climate models. *J. Climate*, 7:914–928, 1994.
- [13] Chris Poulsen, Raymond T. Pierrehumbert, and Robert L. Jacob. Impact of ocean dynamics on the simulation of the neoproterozoic “snowball” earth. *Geophys. Res. Lett.* *in press*, 2001.

- [14] Michael Tobis. *Effect of Slowed Barotropic Dynamics in Parallel Ocean Climate Models*. PhD thesis, University of Wisconsin-Madison, 1996.
- [15] David L. Williamson, Jeffrey T. Kiehl, V. Ramanathan, Robert E. Dickinson, and James J. Hack. Description of NCAR Community Climate Model (CCM1). NCAR Tech. Note NCAR/TN-285+STR, Natl. Cent. for Atmos. Res., Boulder, Co., 1987.